

Abstract of the Invention

- Content is frequently replicated across IP-based networks such as the Internet. Replicas are stored in caches throughout the network. However, to date there has been no mechanism to allow users not directly associated with the replicas to access replicas. Therefore although a copy of a content item may be stored geographically nearby, it may still be necessary to obtain the original content item from a server many thousands of miles away.
- 10 By recording the location of content items and producing a mapping of a content item to all known replicas, it is possible to produce an authoritative proxy associated with a particular domain which is able to provide a list of URLs of alternative replica copies of a requested content item. In this way, a content request proxy associated with a requester of a content item may choose the best location for retrieving the content item
- 15 from a list which may include local caches not directly associated with the requester.